REMARKS

Claims 1 & 3-27 are pending in the application.

Claims 1 & 3-27 have been rejected.

Claim 27 has been amended.

Rejection of Claims under 35 U.S.C. § 101

Claim 27 is rejected under 35 U.S.C. § 101 as reciting software without claiming associated computer hardware. Responsive to this rejection, Applicants have amended Claim 27 to address the Examiner's concern. Specifically, Claim 27 now recites "a network interface means" which is supported, by way of non-limiting example, at paragraph [0060] of the original application.

Rejection of Claims under 35 U.S.C. § 112

Claims 1 and 3-27 are rejected under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which Applicants regard as the invention. Applicants respectfully traverse each of these rejections.

As an initial matter, Applicants respectfully assert that the Office Action's rejections under § 112 are improper, because they relate to claim breadth, rather than indefiniteness. As will be appreciated, if the scope of the claimed subject matter can be determined by one having ordinary skill in the art, a rejection [which refers to indefiniteness] would not be appropriate. *See* MPEP 706.03(d). At page 3, the Office Action asserts:

> The following claim languages are unclear and indefinite:
>
> i) Claim 1, line 3, it is not certain as to what is meant be "subsequent of the cluster being partitioned..." <i.e. was the cluster originally not partitioned and had multiple applications running on the entire cluster? If so, what is the cluster before being partitioned a cluster of, since it is not partitioned? A cluster means a cluster of things or partitions. If it has not been partitioned, how can it be a cluster? The claim lacks a description of how all components, including logical cluster, sub-cluster, and nodes, are related prior to partitioning.>

See Office Action, p.3. The Examiner's concern appears to be that terminology could read on a breadth of multiple embodiments. Breadth of a claim is not to be equated with indefiniteness. *See In re Miller*, 441 F.2d 689, 169 USPQ 597 (CCPA 1971). If the scope of the subject matter embraced by the claims is clear, and if applicants have not otherwise indicated that they intend the invention to be of a scope different from that defined in the claims, then the claims comply with 35 U.S.C. 112, second paragraph. *See* MPEP 2173.04. More specifically, Applicants respectfully submit that questions as to what is meant by common words such as 'subsequent to', the meanings of which are clear and ascertainable to a person of ordinary skill in the art, go to breadth, rather than definiteness. Therefore, Applicants respectfully submit that Claim 1 is definite and request withdrawal of the Office Action's rejections under 35 U.S.C. § 112.

Further, as will be appreciated, in reviewing a claim for compliance with 35 U.S.C. § 112, second paragraph, the examiner must consider the claim as a whole to determine whether the claim apprises one of ordinary skill in the art of its scope and, therefore, serves the notice function required by 35 U.S.C. § 112, second paragraph, by providing clear warning to others as to what constitutes infringement of the patent. *See, e.g., Solomon v. Kimberly-Clark Corp.*, 216 F.3d 1372, 1379, 55 USPQ2d 1279, 1283 (Fed. Cir. 2000). Applicants respectfully submit that, considering the claim as a whole,

the Office Action's question "as to what is meant be 'subsequent of the cluster being partitioned...'" is answered in the following words of the claim limitation. Quoted in full, the claim recites the temporal limitation, "subsequent to the cluster being partitioned into a plurality of sub-clusters." Applicants respectfully submit that the claim language denotes a specifically discernable temporal standard, the fulfillment of which would be apparent to one of ordinary skill in the art. The Office Action asks "was the cluster originally not partitioned and had multiple applications running on the entire cluster?" The Office Action further asks "If so, what is the cluster before being partitioned a cluster of, since it is not partitioned?" Applicants respectfully submit that "the requirement to 'distinctly' claim means that the claim must have a meaning discernible to one of ordinary skill in the art when construed according to correct principles....Only when a claim remains insolubly ambiguous without a discernible meaning after all reasonable attempts at construction must a court declare it indefinite." *See Metabolite Labs., Inc. v. Lab. Corp. of Am. Holdings*, 370 F.3d 1354, 1366, 71 USPQ2d 1081, 1089 (Fed. Cir. 2004). Applicants respectfully submit that "subsequent to the cluster being partitioned into a plurality of sub-clusters" has a readily discernable meaning, and Applicants respectfully request that the Examiner withdraw the Office Action's rejection of claim 1 for indefiniteness.

Similarly, at page 5, the Examiner asserts that, "Claim 8, line 2, it is unclear as to what the relationship is between 'more than one of the subclusters' and 'a plurality of the subclusters' in line 4 of Claim 1." Applicants respectfully submit that the terms "a sub-cluster of the more than one of the sub-clusters continuing to execute a third application of the applications subsequent" and "the cluster being partitioned into a plurality of sub-

clusters" are clear and unambiguous when read in their entirety and in context. Applicants respectfully request that the Examiner withdraw the Office Action's rejection of claim 8 for indefiniteness.

At page 3, the Examiner further asserts:

> It is unclear as to what the relationships is between "first logical cluster" of line 6 and "first one of the sub-clusters" of line 8 <i.e. what does the first logical cluster encompass? does it include the first sub-cluster?>
>
> lines 9-10, it is unclear as to what is meant by 'automatically wins ownership…" <i.e. why would the sub-cluster need to win the ownership of the logical cluster if the logical cluster is already part of the sub-cluster?>

Applicants respectfully refer to the argument advanced above with respect to the difference between breadth and definiteness. Applicants respectfully assert that "why would the sub-cluster need to win the ownership of the logical cluster if the logical cluster is already part of the sub-cluster" is a question of motivation not relevant to determining whether the recited limitation has been violated. Therefore, Applicants respectfully submit that Claim 1 is definite and clear. Applicants respectfully request that the Examiner withdraw the Office Action's rejection of claim 1 for indefiniteness.

As the rejections of Claims 4, 5, 7, 9, 12, 13, 14, 25, 19 and 24 are based on similar questions as to relationships and temporal arrangements, Applicants respectfully request that the Examiner withdraw the Office Action's rejections of Claims 4, 5, 7, 9, 12, 13, 14, 25, 19 and 24 for indefiniteness.

Similarly, the Examiner has complained with respect to Claim 9 that "it is unclear how selecting is done 'based on application-specific information…' <i.e. how does the information determine how the selection is done?>" One having ordinary skill in the art

will be able to ascertain whether the prohibition of the claim is violated without regard to how the violation is accomplished. Applicants respectfully request, on the basis of the argument detailed above, that the Examiner withdraw the Office Action's rejection of Claims 9 for indefiniteness.

The Office Action raises a similar question with respect to "application-specific information" recited in Claim 12 and also asks a question as to motivation. Applicants respectfully refer the Examiner to the arguments provided above with respect to application-specific information and motivation as grounds of rejection. The Office Action further asserts that Claims 18, 23 and 27 "have the same deficiencies as claim 12". Applicants respectfully request, on the basis of the arguments detailed above, that the Examiner withdraw the Office Action's rejection of Claims 12, 18, 23 and 27 for indefiniteness.

Similarly, the Office Action includes a rejection of Claim 15 to the effect that "it is unclear how 'if a second node in another one of the sub-clusters is part of a logical cluster…' in lines 3-5 determines 'whether the node can execute the first one of the applications' in lines 1-2 <i.e. how does one influence the decision?" Applicants respectfully refer the Examiner to the arguments advanced above and request that the Examiner withdraw the Office Action's rejection of Claim 15 for indefiniteness.

Rejection of Claims under 35 U.S.C. § 102

Claims 12-13, 18-19, 23-24, and 27 are rejected under 35 U.S.C. 102(e) as being anticipated by U.S. Publication No. 2003/0187927 by Winchell (*Winchell*). Applicants

respectfully traverse each of these rejections. Independent Claims 12, 18, 23 and 27 include materially the following limitations:

> detecting that a cluster has been partitioned into a plurality of sub-clusters, the cluster executing one or more applications, and one of the plurality of sub-clusters comprising a node; and
> determining whether the node can execute a first one of the applications based on application-specific information associated with the first one of the applications.

See Claim 12. Applicants respectfully submit that *Winchell* fails to provide disclosure of one or more of these claim limitations as they are recited in each of Claims 12, 18, 23 and 27.

Specifically, each of Claims 12, 18, 23 and 27 recites "determining whether the node can execute a first one of the applications based on application-specific information associated with the first one of the applications". The Office Action points to paragraph [0159] and asserts that "since an application keep track of its primary and standby node in case of the failure of the primary node, a node will know if it is allowed to execute that program under either normal conditions or conditions of failure." The cited text of *Winchell* teaches:

> [0159] Let application A provide some service S. Certain modifications to this application can be implemented to make it HA. Assume that two instances of A are run at the same time on two nodes of a cluster. One instance, designated the primary, is actually providing the service S. The other instance, designated the standby, is not providing any service currently, but is ready to assume the role of primary, should the primary instance fail (due to the failure of the node on which the primary is executing). Now, let R be the application state which would be needed by the standby in order to take over the active work of the failed primary. For example, in a cell phone application R might describe the state of all the calls being handled. In general, the standby will get a membership event when the primary fails and it will recover state R from a dataspace and become primary. Specifically, the application will use a dataspace in which it stores (writes) state R and the node id of the primary node. Let

> P_node_id be the name of this latter data item. The dataspace for A, as well as a primary and standby instance of A, is illustrated in FIG. 4.

See Winchell, ¶159. The cited text neither teaches nor suggests, much less discloses the recited "determining whether the node can execute a first one of the applications based on application-specific information associated with the first one of the applications." While the cited text does include "the application state which would be needed by the standby in order to take over the active work of the failed primary", there is no suggestion of "determining whether the node can execute a first one of the applications" based on "the application state" or other application-specific information. Moreover, the Office Action's characterization of the reference, in which it is asserted that "since an application keep track of its primary and standby node in case of the failure of the primary node, a node will know if it is allowed to execute that program under either normal conditions or conditions of failure," suggests that the reference teaches a node-specific determination, rather than the application-specific determination recited in Applicants' Claims 12, 18, 23 and 27.

The Examiner has courteously responded to this argument at Page 16 of the Final Office Action, stating:

> Winchell teaches two instances of an application installed on two separate nodes. In the event of the failure of one instance of the application that is running on a first node, which corresponds to the applicants' application specific information associated with the first one of the applications, where the information is failure of one instance, the other node can be determined to take over and execute that application after failure. If no failure occurs, the first node is determined to continually run the application.

See Final Office Action, Page 16. Applicants respectfully submit that Claim 13 recites "determining whether the node can execute..." Applicants respectfully submit that the

cited text teaches no such determining step. To the contrary, the cited text of *Winchell* begins with the boundary condition "assume that two instances of A are run at the same time on two nodes of a cluster." In *Winchell*, knowledge of failure of the first node, if it used in any determination by *Winchell*, is not used in a determination of "whether the node can execute..." because the cited text of *Winchell* requires the assumption that execution is already in progress.

Further, each of Claims 12, 18, 23 and 27 recites "detecting that a cluster has been partitioned into a plurality of sub-clusters". The Office Action points to paragraph [0180] and asserts that "group of cluster nodes corresponds to sub-clusters." The cited text of *Winchell* teaches:

> [0180] Another example of the system is illustrated in FIGS. 6-11. In this example, the API is configured to designate a first group of cluster nodes 200 to take work items coming in from external sources and put them in a work queue 202. A second group of cluster nodes 204 is designated to process the work items 202. A third group of cluster nodes 206 are designated to handle recovery operations. The third group 206 monitors membership events to facilitate the recovery process.

See Winchell, ¶180. The cited text neither teaches nor suggests, much less discloses the recited "detecting that a cluster has been partitioned into a plurality of sub-clusters." Detection of partitioning is not present in the cited text. At most, the cited text teaches an API assigning tasks to nodes. This mere assignment teaches no detecting of a partition.

The Examiner has courteously responded to this argument at Page 16 of the Final Office Action, stating:

> Winchell teaches registering for membership, where each node needs a membership to be in the cluster (Para 9), It is obvious that the membership information keep track how the cluster has been partitioned

> into nodes. Therefore the cluster knows what nodes or sub-cluster belongs to itself.

See Final Office Action, Page 16. Applicants respectfully submit that no information regarding the recited "detecting that a cluster has been partitioned into a plurality of sub-clusters," is present in the cited text of either paragraph. The newly cited text of *Winchell* teaches:

> [0009] Membership is the set of nodes in the cluster. Applications may register for membership events. A node joining or leaving the cluster results in membership events being delivered. A membership event consists of a set of members and a view number. The view number is an increasing integer value which is agreed upon by the members for a particular membership transition.

See Winchell, ¶9. No information regarding a sub-cluster is presented. Further, Applicants' Claim 12 recites "detecting that a cluster has been partitioned into a plurality of sub-clusters," rather than the "membership information keep track how the cluster has been partitioned into nodes" discussed by the final office action. Applicants respectfully remind the Examiner that "[a]ll words in a claim must be considered in judging the patentability of the claim against the prior art." *In re Wilson*, 424 F.2d 1382, 165 USPQ 494, 496 (CCPA 1970). The assertion that membership in a cluster somehow substitutes for the existence of sub-clusters unfairly parses the claim to eliminate the existence of the recited sub-cluster.

Applicants additionally respectfully submit that the standard for a *prima facie* case of anticipation under 35 U.S.C. § 102 is whether the reference teaches or suggests the features recited in Applicants' claim. Rather than pointing to a teaching of the recited feature within the reference text, the Final Office Action asserts that the recited feature is "obvious." Applicants respectfully submit that the assertion that a feature is "obvious,"

when substituted in place of a teaching from the reference, very nearly constitutes an admission that the reference will not support a *prima facie* case of anticipation under 35 U.S.C. § 102.

For at least these reasons, Applicants respectfully submit that the *Winchell* does not disclose all of the recited limitations of, and therefore does not anticipate, Applicants' Claims 12, 18, 23 and 27, or Claims 13, 19 and 24, which depend therefrom. Applicants therefore respectfully request the Examiner's reconsideration and withdrawal of all rejections as to all claims, and an indication of the allowability of same.

Rejection of Claims under 35 U.S.C. § 103

Claims 1, 3-5 and 8-10 are rejected under 35 U.S.C. 103(a) as being unpatentable over *Eilert* in view of the Examiner's assertion of an obvious modification. In order for a claim to be rendered invalid under 35 U.S.C. §103, the subject matter of the claim as a whole would have to be obvious to a person of ordinary skill in the art at the time the invention was made. *See* 35 U.S.C. §103(a). This requires: (1) the reference(s) must teach or suggest all of the claim limitations; (2) there must be some teaching, suggestion or motivation to combine (or to modify) references either in the references themselves or in the knowledge of the art; and (3) there must be a reasonable expectation of success. *See* MPEP 2143; MPEP 2143.03; *In re Rouffet*, 149 F.3d 1350, 1355-56 (Fed. Cir. 1998).

Claim 1 recites that "if each of the nodes participating in the first logical cluster is included in the first one of the sub-clusters subsequent to the cluster being partitioned, the first one of the sub-clusters automatically wins ownership of the first logical cluster."

The Office Action asserts that, "it would have been obvious to one having ordinary skill in the art at the time of applicant's invention to see, based on Eilert's teaching of claim 1, that if a logical cluser (logical partition), containing its nodes (CPs), is already part of its sub-cluster (partition group), then that sub-cluster would naturally own the logical cluster." Even if *Eilert* were to be construed as teaching or suggesting all of the limitations alleged by the present Office Action, there is still no motivation to modify the reference in the manner suggested by the Office Action. The proper rationales for arriving at a conclusion of obviousness, as suggested by the U.S. Supreme Court in *KSR International Co. v. Teleflex, Inc, et al.*, 127 S. Ct. 1727 (2007), include the following tests for determining a motivation to combine or modify elements from the prior art:

> A. Combining prior art elements according to known methods to yield predictable results;
> B. Simple substitution of one known element for another to obtain predictable results;
> C. Use of a known technique to improve similar devices in the same way;
> D. Applying a known technique to a known device ready for improvement to yield predictable results;
> E. "Obvious to try" – choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success;
> F. Some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention. (emphasis added).

The Office Action uses none of these tests. Rather, the Office Action only offers a conclusory statement that "it would have been obvious to one having ordinary skill in the art at the time of applicant's invention to see, based on Eilert's teaching of claim 1, that if a logical cluser (logical partition), containing its nodes (CPs), is already part of its sub-cluster (partition group), then that sub-cluster would naturally own the logical cluster." This naked assertion of obviousness does not meet the required standard for support and contradicts the teaching of the reference.

Applicants respectfully submit that the Supreme Court in KSR noted that the analysis supporting a rejection under 35 U.S.C. §103 must be made explicit. The Court, quoting *In re Kahn,* stated that "rejections on obviousness cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness." *KSR*, 82 USPQ at 1396. Applicants respectfully submit that "it would have been obvious to one having ordinary skill in the art at the time of applicant's invention to see, based on Eilert's teaching of claim 1, that if a logical cluser (logical partition), containing its nodes (CPs), is already part of its sub-cluster (partition group), then that sub-cluster would naturally own the logical cluster" does not meet this test. Neither this passage, nor the Office Action's statement of motivation, provides evidence of a teaching or some articulated reasoning with some rational underpinning that would motivate a person of skill in the art to modify the reference.

Further, the reference itself deals specifically with failover and teaches directly away from the features articulated in Applicants' claim 2. Rather than deal with failover through the recited, "if each of the nodes participating in the first logical cluster is included in the first one of the sub-clusters subsequent to the cluster being partitioned, the first one of the sub-clusters automatically wins ownership of the first logical cluster," *Eilert* deals with configuration adjustment from Col. 19, line 7 – Column 25, line 40. In the interest of brevity, Applicants decline to quote the text in its entirety. In relevant part, the cited text teaches:

> Returning to FIG. 13a, in addition to the above, the availability index for each path is calculated, STEP 1316. In one example, the availability index is a number that indicates how many single points of failure the proposed path has in common with existing paths to the subsystem. If a channel path

is to be added, then it is desired to have no single points of failure. If a channel path is to be deleted, then the path with the most single points of failure is typically chosen.

Subsequently, the impact to the effected systems is projected, STEP 1318. In particular, in one example, the current load on each subsystem is looked at to determine how it will be different if the change is made. Using this information, the best option is selected, STEP 1320. In order to select the best option, various factors may be considered including, for instance, the following: Which option moves the subsystem closest to target? Which option provides the best availability? Which option provides the best symmetry (least entropy)? Will this option reduce the total number of paths below two? Will this option violate any explicit targets (WLM can provide explicit targets to be used instead of the default targets)? Will this option violate any architectural limits? Will this option violate the configuration as defined by the installation? Will this option attempt to use resources which are currently unavailable?

In one particular example, initially, any decision selection blocks (DSBs) that cannot be implemented under any circumstances are eliminated. This would include, for example, those that violate architectural limits, those that reduce the total number of paths below two, those that violate the configuration as defined by the installation (e.g., uses more than the maximum number of floating channel paths allowed in the definition) and those that attempt to use resources which are currently unavailable (e.g., those that attempt to use ports which were made unavailable.) (This function could be moved earlier in the processing, so that the availability index and projected impact to subsystems is not calculated for DSBs which could never be selected.)

If there is currently only one path to the subsystem (possible shortly after system startup or after a failure), select a path with the best availability index. If more than one have equivalent availability indices, select the path that has the lowest entropy index that will move the target subsystem within the target I/O velocity target. If there are more than one, select the one whose total projected delta (".SIGMA. Projected Delta" from the DSB) is smallest.

If there is currently more than one path to the subsystem, find the set of DSBs with the best availability index. Search that set for options which get the target subsystem within the tolerance of the target I/O velocity and have the lowest entropy index. If there are more than one, select the one who's total projected delta (".SIGMA. Projected Delta" from the DSB) is smallest.

If there is no such option, find the set of paths with the next best availability index, and try again.

If no options get the subsystem within the tolerance, select the one that gets the subsystem closest to target, without regard for its availability index or entropy index. (The above described technique for selecting the best option is only one example. Various additions, deletions and modifications can be made without departing from the spirit of the present invention. Further, any other techniques may be used to select the best option.)

After attempting to select the best option, a determination is made as to whether the new target can be achieved without impacting subsystems with explicit targets, INQUIRY 1322. In other words, is the best option going to negatively impact the subsystems with explicit targets that were set by WLM. If so, then workload manager is invoked to choose the appropriate path, STEP 1324. Specifically, workload manager picks the path and selects new targets for the donor.

Subsequently, or if the new target can be achieved without negatively impacting subsystems with explicit targets, the change is implemented, STEP 1306, and processing continues with a determination as to whether other subsystems exist which are not in target range, INQUIRY 1308.

As stated above, WLM can set an explicit I/O velocity target, which is to be used instead of the default average I/O velocity target. In one embodiment, WLM sets an explicit target, when WLM finds that the service class is not meeting its goals. One embodiment of the logic associated with setting the explicit subsystem I/O velocity target is described with reference to FIG. 15.

Initially, a determination is made as to whether I/O is causing the biggest delay, INQUIRY 1500. If not, then for purposes of this aspect of the invention processing is complete. However, if I/O is causing the biggest delay, then I/O priority adjustment is attempted, STEP 1502. Thereafter, a determination is made as to whether the service class is meeting its goals, INQUIRY 1504. If the service class is now meeting its goals, then processing is complete. However, if the service class is still not meeting its goals, then a search is made for subsystems that are being used by the service class and have a low I/O velocity, STEP 1506. For one or more subsystems that are located, new subsystem I/O velocity targets are set. In one example, it is set by increasing the current target by a defined amount and then projecting the impact on the subsystem. If the impact is

> sufficient, (e.g., above a receiver value), then processing is complete. If not, the target is increased again, and the processing is repeated.
>
> Described in detail above is dynamic CHPID management (DCM) that provides for dynamic adjustment of I/O configurations. DCM is advantageously integrated with WLM, which enables decisions to be made with an understanding of workloads and goals. Further, DCM allows management of channels across multiple partitions (e.g., of a group of partitions). This enables the addition of resources where they are needed, as well as the removal of excess resources.
>
> As described above, with dynamic CHPID management (DCM), the "best" channel is selected to add to (or remove from) a subsystem. In order to do this, one or more attributes are examined, including, for instance, the complexity (or entropy) of the resulting I/O configuration.
>
> Increased entropy causes the I/O configuration to become overly complicated, resulting in excessive time in DCM's processing, inaccurate results due to excessive numbers of subsystems being impacted, complexity in performance reporting and complexity in problem determination. Thus, in one aspect of the present invention, a capability is provided to determine the relative entropy of different choices, so that the relative entropy can be considered along with other considerations, such as I/O velocity and availability, when making a choice in how a configuration is to be adjusted.
>
> In determining relative entropy, an entropy index is calculated. For example, in the case where a path is being added, the entropy index is the sum of channels and subsystems in the resulting configuration. Further, in the case where a path is being deleted, the entropy index reflects the set of channels and subsystems interconnected to the target subsystem after the path is deleted. It does not include channels and subsystems that will no longer be in the set.

See Eilert at Col. 21, line 65 – Col. 23, line 60 (emphasis added).

Applicants apologize for the length of the cited section. The cited text of *Eilert* makes clear, however, that *Eilert* teaches strongly away from the features of Applicants' recited "if each of the nodes participating in the first logical cluster is included in the first one of the sub-clusters subsequent to the cluster being partitioned, the first one of the sub-clusters automatically wins ownership of the first logical cluster." Instead, *Eilert* teaches

to a miasma of entropy indexes, availability indexes, velocity measurements and velocity targets. Applicants respectfully traverse the Examiner's assertion that "it would have been obvious to one having ordinary skill in the art at the time of applicant's invention to see, based on Eilert's teaching of claim 1, that if a logical cluser (logical partition), containing its nodes (CPs), is already part of its sub-cluster (partition group), then that sub-cluster would naturally own the logical cluster." Instead, Applicants respectfully submit that one skilled in the art, upon reading *Eilert*, would be inspired to a solution composed of a complex combination of entropy indexes, availability indexes, velocity measurements and velocity targets, just as *Eilert* teaches. Without any indication of a motivation to modify *Eilert* in a manner that would suggest Applicants' recited limitation, and in the presence of such profoundly contrary teachings within the reference, Applicants respectfully submit that *Eilert* does not render obvious the limitations recited in Applicants' amended Claim 1.

The Examiner has courteously responded to this argument at Page 16 of the Final Office Action, stating:

> Eilert, column 4, lines 12-28 and column 6, lines 29-36, teaches logical partitions (corresponding to applicant's logical cluster) that make up a partition group (corresponding to applicant's sub-cluster). Let:
>
> logical partition (logical cluster) = A
>
> Partition group (sub-cluster) = B
>
> One have, according to above:
>
> A belongs to B
>
> By a simple substitution, one have, the logical partition belongs to sub-cluster, in other words, the sub-cluster owns the logical partition.

See Final Office Action, Page 16. The substitutional logic presented above operates on the implied assumption that, if "B is a part of A, then B owns A." Applicants respectfully

submit that no teaching is offered to support why this should be true. The newly cited text of Column 4 of *Eilert* merely states:

> Computing environment 100 includes, for example, a central processor complex (CPC) 102 having one or more central processors 106 (e.g., CP1-CP4), one or more partitions 108 (e.g., logical partitions (LP1-LP4)), and at least one logical partition manager 110, each of which is described below.
>
> Central processors 106 are physical processor resources that are allocated to the logical partitions. In particular, each logical partition 108 has one or more logical processors (not separately shown for clarity), each of which represents all or a share of a physical processor 106 allocated to the partition. The logical processors of a particular partition 108 may be either dedicated to the partition (so that the underlying processor resource 106 is reserved for that partition) or shared with another partition (so that the underlying processor resource is potentially available to another partition).

See Eilert, column 6, lines 29-36. While the newly cited text addresses availability of a processor to a partition, it does not address the ownership of a cluster by a sub-cluster. A similar deficiency exists in the newly cited text of Column 6 of *Eilert*, which states:

> One embodiment of two logical partition groups (or clusters) of a central processor complex is depicted in FIG. 3. As shown, there is a Logical Partition Group A 300 and a Logical Partition Group B 302, each of which includes one or more logical partitions. The grouping of logical partitions enables resource sharing among the partitions of a group through resource allocation (e.g., priority based resource allocation).

See Eilert, column 6, lines 29-36. Again, there is no mention of the recited "if each of the nodes participating in the first logical cluster is included in the first one of the sub-clusters subsequent to the cluster being partitioned, the first one of the sub-clusters automatically wins ownership of the first logical cluster."

For at least these reasons, Applicants submit that *Eilert* fails to render obvious all of the limitations of independent Claim 1, as amended and that this claim and all claims depending therefrom are in condition for allowance.

Claims 6, 7 and 11 are rejected under 35 U.S.C. § 103 as being unpatentable over *Eilert* in view of *Winchell*. As these claims depend from and further patentably distinguish independent Claim 1, Applicants respectfully request the Examiner's reconsideration and withdrawal of all rejections as to all claims, and an indication of the allowability of same.

Claims 14-17, 20-22 and 25-26 are rejected under 35 U.S.C. § 103 as being unpatentable over *Winchell*. As these claims depend from and further patentably distinguish independent Claims 12, 18 and 23, respectively, Applicants respectfully request the Examiner's reconsideration and withdrawal of all rejections as to all claims, and an indication of the allowability of same.

CONCLUSION

In view of the amendments and remarks set forth herein, the application and the claims therein are believed to be in condition for allowance without any further examination and a notice to that effect is solicited. Nonetheless, should any issues remain that might be subject to resolution through a telephonic interview, the Examiner is invited to telephone the undersigned.

If any extensions of time under 37 C.F.R. § 1.136(a) are required in order for this submission to be considered timely, Applicant hereby petitions for such extensions. Applicant also hereby authorizes that any fees due for such extensions or any other fee associated with this submission, as specified in 37 C.F.R. § 1.16 or § 1.17, be charged to Deposit Account 502306.

Respectfully submitted,

Brenna A. Brock
Attorney for Applicants
Reg. No. 48,509
Telephone: (512) 439-5087
Facsimile: (512) 439-5099